arXiv:1612.06067v1 [math.OC] 19 Dec 2016

A Convex Program for Mixed Linear Regression with a Recovery Guarantee for Well-Separated Data

Paul Hand* and Babhru Joshi*

January 1, 2017

Abstract

We introduce a convex approach for mixed linear regression over d features. This approach is a second-order cone program, based on L1 minimization, which assigns an estimate regression coefficient in $\mathbb{R}^d$ for each data point. These estimates can then be clustered using, for example, k-means. For problems with two or more mixture classes, we prove that the convex program exactly recovers all of the mixture components in the noiseless setting under technical conditions that include a well-separation assumption on the data. Under these assumptions, recovery is possible if each class has at least d independent measurements. We also explore an iteratively reweighted least squares implementation of this method on real and synthetic data.

1 Introduction

We study the mixed linear regression problem with k classes and m data points. Let $\{S_p\}_{p=1}^k$ be a partition of $[m]$, where $i \in S_p$ means that the ith data point belongs to the pth class. In class p, we assume that the data is generated by a linear process with regression coefficients $\beta_p \in \mathbb{R}^d$. Let the measurements $\{(a_i, b_i)\}_{i=1}^m \subset \mathbb{R}^d \times \mathbb{R}$ be such that for all $p \in [k]$,

$$a_i^\top \beta_p = b_i \text{ if } i \in S_p. \tag{1}$$

The mixed linear regression problem is to simultaneously estimate $\{\beta_p\}_{p=1}^k$ and $\{S_p\}_{p=1}^k$ from $\{(a_i, b_i)\}_{i=1}^m$. In the above problem formulation, let $n_p = |S_p|$. We define ℓ_i as the p such that $i \in S_p$, i.e. ℓ_i is the label of the ith data point.

Mixed linear regression has broad applications in scenarios that require disentangling data into multiple latent classes and estimating the parameters. One application is resource management in health care, where the classes are patients (with and without illness) and the mixed data is usage of resource and medical care [5]. Other applications include music tone perception [4], subspace clustering [6] and population classification [7].

Recent work on mixed linear regression has mostly focused on mixtures of two classes, i.e. $k = 2$ [10, 3]. For independent, Gaussian measurements, the authors in [10]

*Department of Computational and Applied Mathematics, Rice University, TX.

develop an initialization procedure for the Expectation Minimization (EM) algorithm. This initialization is based on a grid search on the eigenspace of a matrix of second order moments of measurements. In the noiseless setting, the authors prove that their initialization, followed by an EM algorithm recovers the two regression coefficients with high probability if there are $O(d \log^2 d)$ Gaussian measurements. In [3], the authors lift the mixed linear regression problem to a low rank non-symmetric matrix completion problem. In the noiseless case with $O(d)$ sub-gaussian measurements, the program exactly recovers the two mixture components with high probability. In both of these approaches, the number of mixture components were restricted to exactly two.

For mixed linear regression with two or more mixture components, tensor-based methods have been introduced in [2] and, very recently, in [11]. In [2], low-rank tensor regression and tensor factorization is used to recover the mixture components. In the noiseless setting, their recovery theorem does not guarantee exact recovery for any m. Instead, it establishes that the recovery error is $O(\frac{1}{\sqrt{m}})$. In [11], the method of moments is used to generate an initialization for the EM algorithm. For i.i.d. Gaussian measurements, the algorithm provably recovers the mixture components if the number of measurements is $O(k^{10}d)$. This approach has an optimal sample complexity in the number of feature elements, d, but its scaling with respect to the number of classes, k, is suboptimal. Another drawback of these tensor-based methods is that they are computationally expensive because they operate in high-dimensional spaces.

The work in the present paper brings ideas from convex clustering to the problem of mixed linear regression. The algorithm we introduce is inspired by the algorithm in [8]. In that algorithm, a ℓ_1 minimization in the form of a fused lasso [9] acts to cluster noisy points in $\mathbb{R}^d$. The problem of mixed linear regressions can be viewed as a challenging generalization of clustering in $\mathbb{R}^d$ where every measurement is subject to a codiemsion-1 ambiguity. We note in particular that in the noiseless case, $\mathbb{R}^d$ clustering is trivial whereas mixed linear regression is not. Consequently, this paper focuses only on the case of noiseless measurements in order to grapple with the difficulty of this codimension-1 ambiguity.

The contribution of the present paper is to introduce a convex algorithm for the mixed linear regression problem with two or more components. In the noiseless setting and under some technical conditions that include well-separation of data, this algorithm exactly recovers the mixture components with kd independent measurements. The convex algorithm is based on an ℓ_1 minimization that assigns an estimate regression coefficient in $\mathbb{R}^d$ for each data point. These estimates can then be clustered by, for example, k-means. An estimate for each class can then be found by running k separate regressions.

1.1 Convex Program and Main Result

We introduce a two-step approach for mixed linear regression. This approach introduces a free variable for each data point that acts as the estimate of the mixture component corresponding to that data point. The first step is to solve a second-order cone program to obtain estimated mixture components for each data point, and the second step is to cluster these estimates using k-means. The details of these steps are:

1. Solve the following second order cone program:

$$\begin{aligned} &\underset{z_1,\ldots,z_m\in\mathbb{R}^d}{\arg\min} \sum_{i=1}^{m}\sum_{j=1}^{m} \|z_i - z_j\|_2 \\ &\text{subject to } a_i^\top z_i = b_i,\ i \in [m]. \end{aligned} \tag{2}$$

   Each z_i in (2) is constrained to belong to the hyperplane corresponding to the measurement (a_i, b_i), as per (1). If $\{z_i^\natural\}_{i=1}^m$ is the minimizer of (2), then $z_i^\natural$ is an estimate of β_p for the class p such that $i \in S_p$.

2. Cluster $\{z_i^\natural\}_{i=1}^m$ using k-means and estimate the mixture components corresponding to each class by running k separate regressions.

Program (2) is an ℓ_1 minimization over all pairs (i, j). Due to the sparsity promoting property of ℓ_1 minimization, in some cases, (2) can find a solution where $z_i^\natural = z_j^\natural$ for many pairs (i, j). Let $\{z_i^\sharp\}_{i=1}^m$ be the candidate solution of (2), i.e. $z_i^\sharp = \beta_{\ell_i}$. In the noiseless case, successful recovery would mean that the minimizer of (2) is equal to the candidate solution, i.e. $z_i^\natural = \beta_{\ell_i}$. Even in this case, note that $\|z_i^\natural - z_j^\natural\|_2 \neq 0$ for most pairs (i, j). Nonetheless, we will prove that if the measurements are well-separated in a certain sense, the minimizer of (2) can still be the candidate solution.

We prove that under two technical assumptions, (2) can exactly recover the regression coefficients in the noiseless case. In order to state these assumptions, let

$$v_{pq} = \frac{\beta_p - \beta_q}{\|\beta_p - \beta_q\|_2} \tag{3}$$

and define the weighted directions, v_p, by

$$v_p := \frac{1}{\sum_{q\neq p} n_q} \sum_{q\neq p} n_q v_{pq}. \tag{4}$$

That is, v_p is the weighted average of the directions to β_p from other mixture components $\{\beta_q\}_{q\neq p}$.

Our first assumption is that the measurements are "well-separated" in the following sense:

$$\max_{p\in[k]} \max_{i\in S_p} \frac{\|P_{v_p^\perp} a_i\|_2}{\|P_{v_p} a_i\|_2} < \frac{1}{2} \min_{p\in[k]} \frac{n_p}{m}, \tag{5}$$

where P_{v_p} is the projector onto the span of $\{v_p\}$ and $P_{v_p^\perp}$ is the projector onto the subspace orthogonal to v_p. Intuitively, $\{(a_i, b_i)\}_{i=1}^m$ is well-separated if a_i is approximately parallel to v_{l_i}. This condition ensures that the hyperplanes corresponding to a fixed class are within a small angle of each other. Note that $\frac{\|P_v^\perp a\|_2}{\|P_v a\|_2} = 0$ if v is parallel to a and is small if v is approximately parallel to a. Examples of well-separated data and not well-separated data are shown in Figures 1a and 1b.

Our second assumption is that the measurements are "balanced" in the following sense:

$$\sum_{i\in S_p} \operatorname{sign}(v_p^\top a_i) \frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} = 0, \text{ for all } p \in [k]. \tag{6}$$

Intuitively, $\{(a_i, b_i)\}_{i=1}^m$ is balanced if a particular average of the $\{a_i\}_{i \in S_p}$ is exactly in the direction of v_p. Examples of balanced data and not balanced data are shown in Figures 1a and 1b.

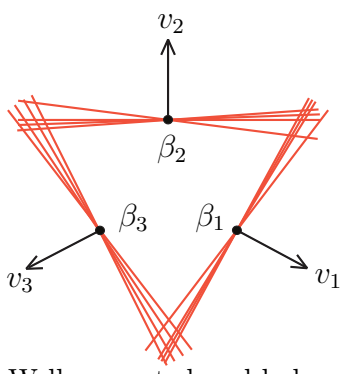

(a) Well-separated and balanced

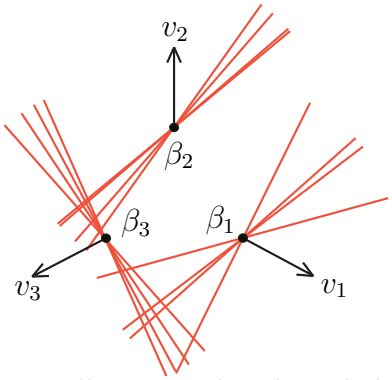

(b) Not well-separated and not balanced

Figure 1: The collection of lines in panel (a) is well-separated in the sense of (5). These measurements also satisfy the balance condition (6). In contrast, the collection of lines in panel (b) is not well-separated because of measurements corresponding to β_1 and is not balanced because of measurements corresponding to β_2. Note that v_p is the weighted average of the directions to β_p from other mixture components $\{\beta_q\}_{q \neq p}$.

Our main result is that in the noiseless case, (2) uniquely recovers the mixture components with well-separated and balanced measurements, provided there are d independent measurements per class.

Theorem 1. *Fix $\{\beta_p\}_{p=1}^k \subset \mathbb{R}^d$. Let $\{(a_i, b_i)\}_{i=1}^m \subset \mathbb{R}^d \times \mathbb{R}$ be measurements that satisfy (1). If the measurements satisfy the well-separation condition* (5)*, the balance condition* (6) *and for all $p \in [k]$, span$\{a_i : i \in S_p\} = \mathbb{R}^d$, then the solution to* (2) *is unique and satisfies $z_i = \beta_{l_i}$ for all $i \in [m]$.*

Thus, mixed linear regression for any number of mixed classes can be exactly solved via a convex program in a case of well-separated data.

The remainder of the paper is organized as follows. In Section 1.2, we present notations used throughout the paper. In Section 2, we present the proof of Theorem 1. In Section 3, we introduce an iteratively reweighted least squares implementation that solves (2). We observe its performance on real and synthetic data.

1.2 Notation

Let $[n] = \{1, \ldots, n\}$. Let $I_{d \times d}$ be the $d \times d$ identity matrix. Let B_r^d be the d dimensional ℓ_2 ball of radius r centered at the origin. For any matrix X and any vector v, let $\|X\|_F$ be the Frobenius norm of X and let $\|v\|_2$ be the ℓ_2 norm of v. For any non-zero vector v, let $\hat{v} = \frac{v}{\|v\|_2}$. For a vector $v \in \mathbb{R}^d$, let $P_{v^\perp} = I_{d \times d} - \hat{v}\hat{v}^\top$ and $P_v = \hat{v}\hat{v}^\top$. For a vector $a_i \in \mathbb{R}^d$, let a_{is} be the sth element of a_i. Throughout the paper, indices i and j are related to data points and indices p and q are related to mixture components.

2 Proof

We will prove Theorem 1 by constructing an exact dual certificate. To arrive at the form of the dual certificate we now derive the KKT conditions for (2). Let $f : (\mathbb{R}^d)^m \to \mathbb{R}$ be

$$f(Z) = f(z_1, \ldots, z_m) = \sum_{i=1}^{m} \sum_{j=1}^{m} \|z_i - z_j\|_2.$$

The augmented Lagrangian for (2) is

$$\mathcal{L}(Z, \nu) = \sum_{i=1}^{m} \sum_{j=1}^{m} \|z_i - z_j\|_2 - \sum_{i=1}^{m} \nu_i \left(a_i^\top z_i - b_i \right),$$

where $\nu \in \mathbb{R}^m$ is a Lagrange multiplier. Let $Z^\sharp$ be the candidate solution, i.e. $z_i^\sharp = \beta_{l_i}$ for all $i \in [m]$. The first order optimality conditions for (2) at $Z^\sharp$ are

$$0 \in \partial_Z \left(\sum_{i=1}^{m} \sum_{j=1}^{m} \|z_i - z_j\|_2 \right) (Z^\sharp) - \partial_Z \left(\sum_{i=1}^{m} \nu_i \left(a_i^\top z_i - b_i \right) \right) (Z^\sharp), \tag{7}$$

$$a_i{}^T z_i^\sharp = b_i, \ i \in [m], \tag{8}$$

where $\partial_Z f(Z^\sharp)$ is the subdifferential of f with respect to Z evaluated at $Z^\sharp$. Note that

$$\partial_{(z_i, z_j)} \|z_i - z_j\|_2 (z_i^\sharp, z_j^\sharp) = \begin{cases} \left\{ \begin{bmatrix} v_{pq} \\ -v_{pq} \end{bmatrix} \right\}, & \ell_i = p, \ \ell_j = q, \ p \neq q \\ \left\{ \begin{bmatrix} \xi_{ij} \\ -\xi_{ij} \end{bmatrix} : \|\xi_{ij}\|_2 \leq 1 \right\}, & z_i^\sharp = z_j^\sharp \end{cases} \tag{9}$$

where v_{pq} is defined in (3) and $\xi_{ij} \in \mathbb{R}^d$. So, the KKT conditions for (2) are: for all $p \in [k]$ and $i \in S_p$,

$$0 \in \left\{ \sum_{j \in S_p, j \neq i} \xi_{ij} + \sum_{q \neq p} n_q v_{pq} - \nu_i a_i : \xi_{ij} \in \mathbb{R}^d, \ \|\xi_{ij}\|_2 \leq 1, \ \xi_{ij} = -\xi_{ji}, \ \nu_i \in \mathbb{R} \right\} \tag{10}$$

We will now prove that if such ξ_{ij} and ν exist and if all $\|\xi_{ij}\|_2 < 1$, then the unique output of (2) is the candidate solution $Z^\sharp$.

Lemma 1. *Fix $\{\beta_p\}_{p=1}^k \subset \mathbb{R}^d$. Let $\{(a_i, b_i)\}_{i=1}^m \subset \mathbb{R}^d \times \mathbb{R}$ be measurements that satisfy (1). Assume for $p \in [k]$ and for all $i, j \in S_p$, there exists $\xi_{ij} \in \mathbb{R}^d$ and $\nu_i \in \mathbb{R}$ such that*

$$\nu_i a_i = \sum_{j \in S_p, j \neq i} \xi_{ij} + \sum_{q \neq p} n_q v_p, \tag{11}$$

$$\|\xi_{ij}\|_2 < 1, \tag{12}$$

$$\xi_{ij} = -\xi_{ji}. \tag{13}$$

Also, assume for all $p \in [k]$, span$\{a_i : i \in S_p\} = \mathbb{R}^d$. Then, $Z^\sharp$ is the unique solution to (2).

Proof. We will show that $f(Z^\sharp + H) > f(Z^\sharp)$ for any feasible perturbation $H \neq 0$. We first separate $f(Z^\sharp + H)$ into two parts.

$$f(Z^\sharp + H) = \sum_{p=1}^{k} \sum_{i,j \in S_p} \|h_i - h_j\|_2 + \underbrace{\sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} \|z_i^\sharp - z_j^\sharp + h_i - h_j\|_2}_{I} \tag{14}$$

To bound I from below, note that $\|z_i^\sharp - z_j^\sharp + h_i - h_j\|_2 \geq \|z_i^\sharp - z_j^\sharp\|_2 + v_{pq}^\top (h_i - h_j)$ because v_{pq}, defined in (3), is a subgradient of $\|\cdot\|_2$ at $z_i^\sharp - z_j^\sharp$ if $i \in S_p$ and $j \in S_q$. Thus,

$$\begin{aligned}
I =& \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} \|z_i^\sharp - z_j^\sharp + h_i - h_j\|_2 && (15)\\
\geq& \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} \|z_i^\sharp - z_j^\sharp\|_2 + \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} v_{pq}^\top (h_i - h_j) && (16)\\
=& f(Z^\sharp) + \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} v_{pq}^\top h_i - \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} \sum_{j \in S_q} v_{pq}^\top h_j && (17)\\
=& f(Z^\sharp) + \sum_{p=1}^{k} \sum_{q \neq p} \sum_{i \in S_p} n_q v_{pq}^\top h_i + \sum_{p=1}^{k} \sum_{q \neq p} \sum_{j \in S_q} n_p v_{qp}^\top h_j && (18)\\
=& f(Z^\sharp) + 2 \sum_{p=1}^{k} \sum_{i \in S_p} \sum_{q \neq p} n_q v_p^\top h_i. && (19)
\end{aligned}$$

Note that in (18) we used $v_{pq} = -v_{qp}$ and in (19) we used (4). Combining (14) and (19), it suffices to show

$$\sum_{p=1}^{k} \sum_{i,j \in S_p} \|h_i - h_j\|_2 + 2 \underbrace{\sum_{p=1}^{k} \sum_{i \in S_p} \sum_{q \neq p} n_q v_p^\top h_i}_{II} > 0. \tag{20}$$

for all feasible $H \neq 0$.

First, we show that for all feasible $H \neq 0$, there exists $p \in [k]$ and $i, j \in S_p$ such that $h_i \neq h_j$. We provide a proof by contradiction. Assume for all $p \in [k]$ and $i, j \in S_p$, $h_i = h_j$. Fix $p \in [k]$. Let $h_j = c_p$ for all $j \in S_p$. Since c_p is a feasible perturbation, $a_j^\top c_p = 0$ for all $j \in S_p$. Thus, c_p is orthogonal to any element in $\text{span}\{a_j : j \in S_p\} = \mathbb{R}^d$. Hence, $h_j = c_p = 0$ for all $j \in S_p$, which contradicts $H \neq 0$.

We now compute

$$\begin{aligned}
II =& \sum_{p=1}^{k} \sum_{i \in S_p} \left(\sum_{q \neq p} n_q v_p \right)^\top h_i && (21)\\
=& \sum_{p=1}^{k} \sum_{i \in S_p} \left(\nu_i a_i - \sum_{j \in S_p, j \neq i} \xi_{ij} \right)^\top h_i && (22)
\end{aligned}$$

$$= -\frac{1}{2}\sum_{p=1}^{k} 2 \sum_{i\in S_p} \sum_{j\in S_p, j\neq i} \xi_{ij}^\top h_i \tag{23}$$

$$= -\frac{1}{2}\sum_{p=1}^{k} \left(\sum_{i\in S_p} \sum_{j\in S_p, j\neq i} \xi_{ij}^\top h_i + \sum_{j\in S_p} \sum_{i\in S_p, i\neq j} \xi_{ji}^\top h_j \right) \tag{24}$$

$$= -\frac{1}{2}\sum_{p=1}^{k} \sum_{i\in S_p} \sum_{j\in S_p, j\neq i} \xi_{ij}^\top (h_i - h_j) \tag{25}$$

$$\geq -\frac{1}{2}\sum_{p=1}^{k} \sum_{i\in S_p} \sum_{j\in S_p} \|\xi_{ij}\|_2 \|h_i - h_j\|_2 \tag{26}$$

$$= -\frac{\gamma}{2}\sum_{p=1}^{k} \sum_{i\in S_p} \sum_{j\in S_p} \|h_i - h_j\|_2, \tag{27}$$

where $\gamma := \max_{p\in[k]} \max_{i,j\in S_j} \|\xi_{ij}\|_2$. By assumption (12), $\gamma < 1$. Note that in (22), we used (11). In (23), we used $a_i^\top h_i = 0$ (since H is feasible). In (24), both terms in parenthesis are equal by interchanging the dummy variables i and j. In (25), we used the antisymmetry condition (13). In (26), we used Cauchy-Schwartz inequality. Combining (20) and (27), we get

$$\sum_{p=1}^{k} \left(\sum_{i,j\in S_p} \|h_i - h_j\|_2 + 2 \sum_{i\in S_p} \sum_{q\neq p} n_q v_p^\top h_i \right) \geq (1-\gamma) \sum_{p=1}^{k} \sum_{i,j\in S_p} \|h_p - h_q\|_2 > 0. \tag{28}$$

The strict inequality in (28) holds because $(1-\gamma) > 0$ and $H \neq 0$ implies that there exists $p \in [k]$ and $i, j \in S_p$ such that $h_i \neq h_j$. Hence, $Z^\sharp$ is the unique solution to (2) □

As a consequence of Lemma 1, the proof of Theorem 1 is simplified to constructing an exact certificate that satisfies (11), (12) and (13).

Proof of Theorem 1. For all $p \in [k]$ and $i, j \in S_p$, let ξ_{ij} and ν_i be

$$\nu_i = \operatorname{sign}(v_p^\top a_i) \frac{\|v_p\|_2 \sum_{q\neq p} n_q}{\|P_{v_p} a_i\|_2} \tag{29}$$

$$\xi_{ij} = \frac{1}{n_p} \left(\nu_i P_{v_p^\perp} a_i - \nu_j P_{v_p^\perp} a_j \right). \tag{30}$$

By Lemma 1, it is sufficient to show that, for all $p \in [k]$ and $i, j \in S_p$, these ξ_{ij} and ν_i satisfy (11), (12), (13). Note that (13) follows immediately. Condition (12) holds because

$$\|\xi_{ij}\|_2 \leq \frac{1}{n_p} \left(\|\nu_i P_{v_p^\perp} a_i\|_2 + \|\nu_j P_{v_p^\perp} a_j\|_2 \right) \tag{31}$$

$$\leq \frac{2}{n_p} \max_{i\in S_p} \|\nu_i P_{v_p^\perp} a_i\|_2 \tag{32}$$

$$= \frac{2}{n_p} \max_{i\in S_p} \|v_p\|_2 \sum_{q\neq p} n_q \frac{\|P_{v_p^\perp} a_i\|_2}{\|P_{v_p} a_i\|_2} \tag{33}$$

$$< \frac{2}{n_p} m \frac{n_p}{2m} \tag{34}$$

$$= 1 \tag{35}$$

In (33), we used (29), and in (34), we used (5) along with $\|v_p\|_2 \leq 1$.

Lastly, we will show that (11) holds through direct computation. Fix $p \in [k]$ and $i \in S_p$. We now compute

$$\sum_{j \in S_p, j \neq i} \xi_{ij} + \sum_{q \neq p} n_q v_p \tag{36}$$

$$= \sum_{j \in S_p, j \neq i} \frac{1}{n_p} \left(\nu_i P_{v_p^\perp} a_i - \nu_j P_{v_p^\perp} a_j \right) + \sum_{q \neq p} n_q v_p \tag{37}$$

$$= \sum_{j \in S_p, j \neq i} \frac{1}{n_p} \|v_p\|_2 \left(\sum_{q \neq p} n_q \right) \left(\operatorname{sign}(v_p^\top a_i) \frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} - \operatorname{sign}(v_p^\top a_j) \frac{P_{v_p^\perp} a_j}{\|P_{v_p} a_j\|_2} \right) + \sum_{q \neq p} n_q v_p \tag{38}$$

$$= \left(\sum_{q \neq p} n_q \right) \frac{\|v_p\|_2}{n_p} \sum_{j \in S_p} \left(\operatorname{sign}(v_p^\top a_i) \frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} - \operatorname{sign}(v_p^\top a_j) \frac{P_{v_p^\perp} a_j}{\|P_{v_p} a_j\|_2} \right) + \sum_{q \neq p} n_q v_p \tag{39}$$

$$= \left(\sum_{q \neq p} n_q \right) \frac{\|v_p\|_2}{n_p} \left(\operatorname{sign}(v_p^\top a_i) n_p \frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} - \sum_{j \in S_p} \operatorname{sign}(v_p^\top a_j) \frac{P_{v_p^\perp} a_j}{\|P_{v_p} a_j\|_2} \right) + \left(\sum_{q \neq p} n_q \right) v_p \tag{40}$$

$$= \left(\sum_{q \neq p} n_q \right) \|v_p\|_2 \left(\operatorname{sign}(v_p^\top a_i) \frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} + \frac{v_p}{\|v_p\|_2} \right) \tag{41}$$

$$= \left(\sum_{q \neq p} n_q \right) \frac{\|v_p\|_2 \operatorname{sign}(v_p^\top a_i)}{\|P_{v_p} a_i\|_2} \left(P_{v_p^\perp} a_i + \operatorname{sign}(v_p^\top a_i) \frac{v_p}{\|v_p\|_2} \|P_{v_p} a_i\|_2 \right) \tag{42}$$

$$= \nu_i \left(P_{v_p^\perp} a_i + P_{v_p} a_i \right) \tag{43}$$

$$= \nu_i a_i. \tag{44}$$

Note that in (38), we used (29). In (41), we used the assumed balance condition (6), and in (43), we used (29) and $P_{v_p} a_i = \operatorname{sign}(v_p^\top a_i) \|P_{v_p} a_i\|_2 \frac{v_p}{\|v_p\|_2}$. Hence, (11) holds. By Lemma 1, $Z^\sharp$ is the unique solution to (2). □

2.1 Derivation of Dual Certificate

In this section, we provide a derivation of constructing the dual certificate. Fix $p \in [k]$. Without loss of generality, assume $S_p = [n_p]$.

Note that (11) can be decomposed into its components along and orthogonal to v_p:

$$\sum_{j \in S_p, j \neq i} P_{v_p} \xi_{ij} + \sum_{q \neq p} n_q v_p = \nu_i P_{v_p} a_i, \ i \in S_p \tag{45}$$

$$\sum_{j\in S_p, j\neq i} P_{v_p^\perp}\xi_{ij} = \nu_i P_{v_p^\perp} a_i,\ i \in S_p. \tag{46}$$

Note that solving (46) is equivalent to solving a system of linear equations $A\xi = b$, where $A \in \mathbb{R}^{n_p d \times \frac{n_p(n_p-1)d}{2}}$ and $b \in \mathbb{R}^{n_p d}$. Let $\tilde{B}$ be the first (n_p-1) block rows of A and $\tilde{b}$ be the first (n_p-1) blocks of b. So, the ith block of $\tilde{b}$ is $\nu_i P_{v_p^\perp} a_i$ and the matrix $\tilde{B}$ and the vector ξ are

$$(\tilde{B})_{il} = \begin{cases} I_{d\times d} & \text{if } l \in \left\{ \sum_{r=1}^{i-1}(n_p - r) + s :\ s = 1, \ldots, n_p - i \right\} \\ -I_{d\times d} & \text{if } i > 1,\ l \in \left\{ \sum_{r=1}^{i-s}(n_p - r) + i - s + 1 :\ s = 2, \ldots, i \right\} \\ 0 & \text{otherwise,} \end{cases}$$

$$(\xi)_l = P_{v_p^\perp}\xi_{ij} \text{ if } i = \underset{\tilde{i}\in[n_p-1]}{\arg\max} \sum_{s=1}^{\tilde{i}} \frac{l - \sum_{r=1}^{s-1}(n_p - r)}{\left| l - \sum_{r=1}^{s-1}(n_p - r) \right|},\ j = l - \sum_{r=1}^{i-1}(n_p - r) + i.$$

It is straightforward to verify that $\tilde{B}$ is full row rank and the least squares solution of $\tilde{B}\xi = \tilde{b}$ is also a solution to $A\xi = b$ if

$$\sum_{i=1}^{n_p} \nu_i P_{v_p^\perp} a_i = 0. \tag{47}$$

Now, for all $j \in S_p$, choose $P_{v_p}\xi_{ij} = 0$ in (45). Then, for all $i \in S_p$,

$$\nu_i = \text{sign}(v_p^\top a_i) \frac{\|v_p\|_2 \sum_{q\neq p} n_q}{\|P_{v_p} a_i\|_2}. \tag{48}$$

Using these ν_i, (47) is satisfied because of the assumed balance condition (6). We note that the least 2-norm solution of $\tilde{B}\xi = \tilde{b}$ is $\xi = \tilde{B}^T(\tilde{B}\tilde{B}^T)^{-1}\tilde{b}$, where

$$(\tilde{B}\tilde{B}^T)_{ls} = \begin{cases} (n_p - 1)I_{d\times d} & \text{if } l = s \\ -I_{d\times d} & \text{if } l \neq s, \end{cases} \quad \text{and}$$

$$((\tilde{B}\tilde{B}^T)^{-1})_{ls} = \begin{cases} \frac{2}{n_p} I_{d\times d} & \text{if } l = s \\ \frac{1}{n_p} I_{d\times d} & \text{if } l \neq s. \end{cases}$$

Here, $\tilde{B}\tilde{B}^T \in \mathbb{R}^{d(n_j-1)\times d(n_j-1)}$. Solving for ξ, we get $P_{v_p^\perp}\xi_{ij} = \frac{\nu_i P_{v_p^\perp} a_i}{n_p} - \frac{\nu_j P_{v_p^\perp} a_j}{n_p}$. Since $P_{v_p}\xi_{ij} = 0$, we have $\xi_{ij} = \frac{\nu_i P_{v_p^\perp} a_i}{n_p} - \frac{\nu_j P_{v_p^\perp} a_j}{n_p}$.

3 Numerical Results

In this section we formulate an iteratively reweighted least squares (IRLS) algorithm for (2) and provide numerical results on real and synthetic data. The pseudocode is given in Algorithm 1. We refer the readers to [1] for convergence result of the IRLS algorithm.

Algorithm 1 outputs estimated mixture components for each data point. These

estimates can then be clustered using k-means. The mixture component corresponding to each class is obtained by running k separate regressions.

When implementing the IRLS algorithm, we observe convergence for a fixed $\delta \ll 1$ (δ is fixed to 10^{-16}). The maximum number of iterations is 150. Let $Z_t^\natural$ be the minimizer of the quadratic program in the IRLS algorithm at iteration t. The stopping criteria is $\frac{1}{\sqrt{m}}\|Z_{t+1}^\natural - Z_t^\natural\|_F < 10^{-5}$.

Algorithm 1 Iteratively Reweighted Least Squares for (2)

Input: Measurements $\{(a_i, b_i)\}_{i=1}^m$, regularization parameter $\delta_t \searrow 0$, maximum iteration number t_{max}, $w_{ij}^{(0)} = 1$.
Output: Estimates $\{z_1, \ldots, z_m\}$
while $t \leq t_{max}$ or convergence **do**:
$z_1^{(t)}, \ldots, z_m^{(t)} \leftarrow \arg\min_{z_1,\ldots,z_m} \sum_{i=1}^m \sum_{j=1}^m w_{ij}^{(t)} \|z_i - z_j\|_2^2$ s.t. $a_i^\top z_i = b_i,\ i \in [m]$
$w_{ij}^{(t+1)} \leftarrow \left(\|z_i^{(t)} - z_j^{(t)}\|_2^2 + \delta_t\right)^{-\frac{1}{2}}$
end while

We now use Algorithm 1, along with k-means, to classify two real datasets. First, we analyze BikeShare data that contains the number of bikes rented in a day in the San Francisco bay area from September 1, 2014 to August 30, 2015. The data is provided by the Bay Area BikeShare program. In this dataset, $m = 364$ and $d = 2$. For each measurement $(a_i, b_i) \in \mathbb{R}^2 \times \mathbb{R}$, $a_{i1} = i$, $a_{i2} = 1$ and b_i is the number of bikes rented. Let $\mu = \sum_{i=1}^m a_{i1}$ and $\alpha = 0.005$. Let $\tilde{a}_{i1} = \alpha(a_{i1} - \mu)$ and let $\tilde{a}_i = (\tilde{a}_{i1}, a_{i2})$. Algorithm 1 is then used on the dataset $\{(\tilde{a}_i, b_i)\}_{i=1}^m$. K-means with $k = 2$ is used on the output of Algorithm 1 to estimate the mixture components. The result of this process is shown in Figure 2b.

Second, we consider music tone perception data which shows the relationship between actual tone and tone perceived by a musician. This data was generated in an experiment conducted by Cohen in 1980 [4]. In this dataset, $m = 150$ and $d = 2$. For each measurement $(a_i, b_i) \in \mathbb{R}^2 \times \mathbb{R}$, a_{i1} is the actual tone, $a_{i2} = 1$ and b_i is the perceived tone. Let $\mu = \sum_{i=1}^m a_{i1}$ and $\alpha = 10$. Let $\tilde{a}_{i1} = \alpha(a_{i1} - \mu)$ and let $\tilde{a}_i = (\tilde{a}_{i1}, a_{i2})$. Algorithm 1 is then used on the dataset $\{(\tilde{a}_i, b_i)\}_{i=1}^m$. K-means with $k = 2$ is used on the output of Algorithm 1 used to estimate the mixture components. The result of this process is shown in Figure 3b.

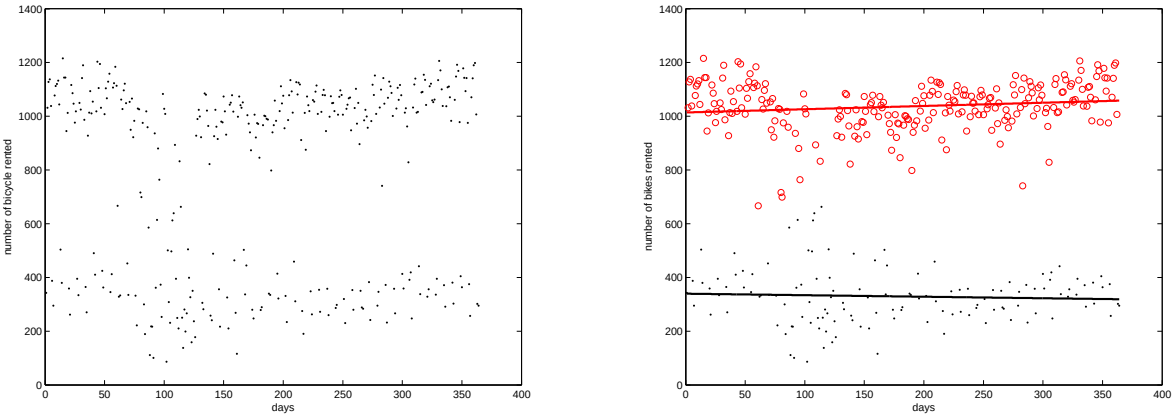

(a) (b)

Figure 2: Panel (a) shows data consisting of the number of bikes rented per day in the Bay Area BikeShare program. Panel (b) shows the results of (2), followed by k-means with $k = 2$. In panel (b), points in the first and second classes are represented by dots and circles, respectively. The lines correspond to estimated mixture components for each class obtained using two separate regressions.

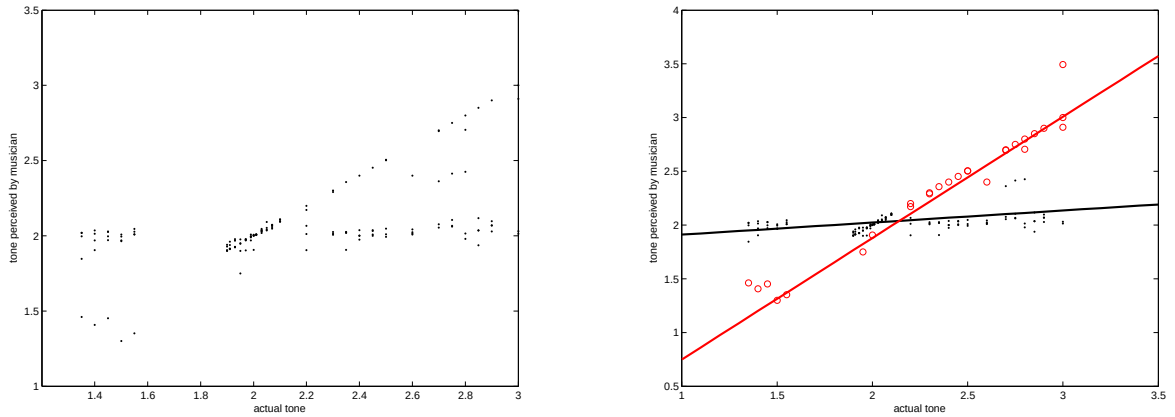

(a) (b)

Figure 3: Panel (a) shows perceived tone by a musician versus actual tone for a range of tones. Panel (b) shows the results of (2), followed by k-means with $k = 2$. In panel (b), points in the first and second classes are represented by dots and circles, respectively. The lines correspond to estimated mixture components for each class obtained using two separate regressions.

We also provide simulation results on synthetic data that support Theorem 1. Let $k = 3$, $d \in \{3, \ldots, 15\}$ and $m = 48$. Let $\{\beta_p\}_{p=1}^k \subset \mathbb{R}^d$ be such that $(\beta_p)_l = 1$ if $l = p$ and 0 otherwise. Let $n_1 = n_2 = n_3 = 16$. For $P_{v_p^\perp}$, let the columns of $Q_p \in \mathbb{R}^{d\times(d-1)}$ be an orthonormal basis of the column space of $P_{v_p^\perp}$.

Consider the following measurements: Fix $p \in [k]$. For $i \in [\frac{n_p}{2}]$, let $x_i \sim \text{Uniform}(B_\alpha^{d-1})$, i.e. x_i is a random point uniformly distributed in the $d-1$ dimensional ℓ_2 ball of radius α centered at the origin. Here, hyperplanes corresponding to labels in S_p are contained in an aperture $\alpha \in [0, 0.75]$. For $i \in [\frac{n_p}{2}]$, let $a_i = \hat{v}_p + Q_p x_i$ and for $i \in \{\frac{n_p}{2}+1, \ldots, n_p\}$, let $a_i = \hat{v}_p - Q_p x_{i-\frac{n_p}{2}}$. These measurements are symmetric and satisfy the balance con-

dition since there exists pairs $i, j \in S_p$ such that

$$\frac{P_{v_p^\perp} a_i}{\|P_{v_p} a_i\|_2} = -\frac{P_{v_p^\perp} a_j}{\|P_{v_p} a_j\|_2}. \tag{49}$$

Lastly, for $i \in [n_p]$, let $b_i = a_i^\top \beta_p$.

Figure 4 shows the fraction of successful recovery from 10 independent trails for mixed linear regression from data as described above. Black squares correspond to no successful recovery and white squares to 100% successful recovery. Let $Z^\sharp$ be the candidate minimizer of (2) and let $Z^\natural$ be the output of (2). For each trail, we say (2) successfully recovers the mixture components if $\frac{1}{\sqrt{m}}\|Z^\natural - Z^\sharp\|_F < 10^{-5}$. In the area to the left of the line, the measurements are well-separated in the sense of (5), i.e. $\max_{p\in[3]} \max_{i\in S_p} \frac{\|P_{v_p^\perp} a_i\|_2}{\|P_{v_p} a_i\|_2} < \frac{1}{6}$. The figure also shows that when measurements are not well-separated, recovery using (2) will likely fail.

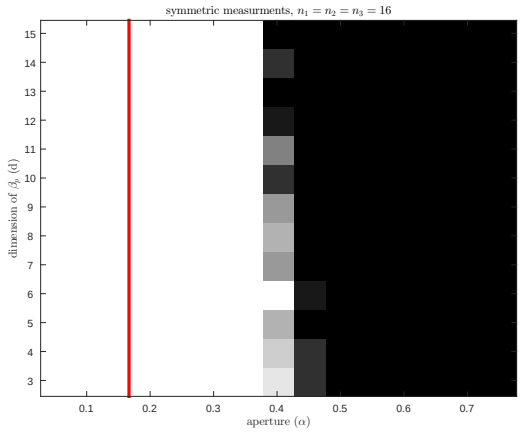

Figure 4: The empirical recovery probability from synthetic data with three mixture components as a function of dimension, d, and aperture, α. The shades of black and white represents the fraction of successful simulation. White blocks correspond to successful recovery and black blocks correspond to unsuccessful recovery. Each block corresponds to the average from 10 independent trials.

References

[1] N. Bissantz, L. Dmbgen, A. Munk, and B. Stratmann. *Convergence analysis of generalized iteratively reweighted least squares algorithms on convex function spaces.* SIAM Journal on Optimization, 19:18281845, 2009.

[2] A. Chaganty and P. Liang. *Spectral experts for estimating mixtures of linear regressions.* International Conference on Machine Learning (ICML), 2013.

[3] Y. Chen, X. Yi, C. Caramanis. *A convex formulation for Mixed Regression with Two Components: Minimax Optimal Rates.* JMLR: Workshop and Conference Proceedings, 35: 145, 2014.

[4] E. Cohen. *Inharmonic tone perception.* Ph.D Dissertation, Stanford University, 1980.

[5] P. Deb and A. M. Holmes. *Estimates of use and costs of behavioural health care: a comparison of standard and finite mixture models*. Health Economics, 9(6): 475-489, 2000.

[6] E. Elhamifar and R. Vidal. *Sparse Subspace Clustering: Algorithm, Theory, and Applications*. 2013.

[7] B. Grun, F. Leisch, et al. *Applications of finite mixtures of regression models*. URL: https://cran.r-project.org/web/packages/flexmix/vignettes/regression-examples.pdf, 2007.

[8] T. Hocking, A. Joulin, F. Bach, J. Vert. *Clusterpath: An Algorithm for Clustering using Convex Fusion Penalties*. ICML, 2011.

[9] R. Tibshirani, M. Saunders, et al. *Sparsity and smoothness via the fused lasso*. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 67: 91-108, 2005.

[10] X. Yi, C. Caramanis, S. Sanghavi. *Alternating Minimization of Mixed Linear Regression*. Jorunal of Machine Learning Research, 2014.

[11] X. Yi, C. Caramanis, S. Sanghavi. *Solving a Mixture of Many Random Linear Equations by Tensor Decomposition and Alternating Minimization*, arXiv preprint arXiv:1608.05749, 2016.